Format: paper
Title: An AI Approach to Ground Station Autonomy for Deep Space Communications
Authors: Forest Fisher·Tara Estlin, Darren Mutz, Leslie Paal, Emily Law, Mike Stockett, Nasser Golshan, Steve Chien
Presenter: Russell Knight
Address: Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Mail Stop 126-347, Pasadena, CA 91109-8099, USA
Phone: (818) 393 5368   Fax: (818) 393 5244
Email: forest.fisher@jpl.nasa.gov
Region: America

This paper describes an architecture for an autonomous deep space tracking station (DS-T). The architecture targets fully automated routine operations encompassing scheduling and resource allocation, antenna and receiver predict generation. track procedure generation from service requests, and closed loop control and error recovery for the station subsystems. This architecture has been validated by the construction of a prototype DS-T station, which has performed a series of demonstrations of autonomous ground station control for downlink services with NASA's Mars Global Surveyor (MGS).

Extended Abstract:
The Deep Space Network (DSN) was established in 1958 and has since evolved into the largest and most sensitive scientific telecommunications and radio navigation network in the world. The purpose of the DSN is to support unmanned interplanetary spacecraft missions and to support radio and radar astronomy observations taken in the exploration of space. The function of the DSN is to receive telemetry signals from spacecraft, transmit commands that control spacecraft operating modes, generate the radio navigation data used to locate and guide a spacecraft to its destination. and acquire flight radio science. radio and radar astronomy, very long baseline interferometry (VLBI), and geodynamics measurements.

This paper will describe the Deep Space Terminal (DS-T), a prototype 34-meter deep space communications station developed as a technology demonstration of fully autonomous, lights-out, operations. In the DS-T concept, a global DSN schedule is disseminated to a set of autonomous DS-T stations. Each DS-T station operates autonomously, performing tracks in a largely independent fashion. When requested to perform a track, the DS-T station performs a number of tasks (at appropriate times) required to execute the track. First, the DS-T station uses appropriate spacecraft navigation ephemeris and predict generation software in order to produce necessary antenna and receiver predict information required to perform the track. Next, the DS-T station executes the pre-calibration process, in which the antenna and appropriate subsystems (e.g., receiver, exciter, telemetry processor, etc.) are configured in anticipation of the track. During the actual track, the signal from the spacecraft must be acquired and the antenna and subsystems must be commanded to retain the signal. adjust for changes in the signal (such as changes in bit rate or modulation index as transmitted by the spacecraft). and perform error recovery Finally, at the completion of the track, the station must be returned to an appropriate standby state in preparation for the next track. All of these activities require significant automation and robust execution including closed loop control. retries and contingency handling.

In order to provide this autonomous operation capability, the DS-T station employs tightly coupled state of the art hardware and software. At the core of the autonomy are two areas of artificial intelligence (AI) technology, AI scheduling and AI planning. We will offer a brief example of each and a brief context for how they apply to the DS-T.

When the decision is made to fly a mission, a very knowledge-intensive process begins that will ensure the necessary DSN antenna coverage. First, a forecast is made of the DSN resources that the spacecraft will require. In the Resource Allocation Process (RAP), the types of services, frequency, and duration of the required tracks are determined as well as high-level resource requirements (e.g., antenna). While the exact timing of the tracks is not known. a set of automated forecasting tools are used to estimate network load and to assist in ensuring that adequate network resources will be available. One part of the network architecture is a unified tool suite that has been developed called TMOD Integrated Ground Resource Allocation System (TIGRAS), which uses operations research and probabilistic reasoning techniques to allow forecasting and capacity planning for DSN resources.

As the time of the actual tracks approaches, this estimate of resource loading is converted to an actual schedule, which becomes more concrete as time progresses. In this process, specific project service requests and priorities are matched up with available resources in order to meet communications needs for earth-orbiting and deep space spacecraft. This scheduling process involves considerations of thousands of possible tracks, tens of projects, tens of antenna resources and considerations of hundreds of subsystem configurations. In addition to adding the detail of antenna subsystem allocation, the initial sched-

ule undergoes continual modification due to changing project needs, equipment availability, and weather considerations. Responding to changing context and minimizing disruption while rescheduling is a key issue.

An evolution of the Operation Mission Planner (OMP-26M) system, the Demand Access Network Scheduler (DANS) is designed to deal with the more complex subsystem and priority schemes required to schedule the larger 34 and 70 meter antennas. Because of the size and complexity of the rescheduling task, manual scheduling is prohibitively expensive. Automation of these scheduling functions is projected to save millions of dollars per year in DSN operations costs.

DANS uses priority-driven, best-first, constraint-based search and iterative optimization techniques to perform priority-based rescheduling in response to changing network demand. In these techniques, DANS first considers the antenna allocation process, as antennas are the central focus of resource contention. After establishing a range of antenna options, DANS then considers allocation of the 5-13 subsystems per track (out of the tens of shared subsystems at each antenna complex) used by each track.

This schedule is first used at a network wide level designating what resources (primarily the antennas) shall be used to provide what services (primarily communications tracks). In the DS-T architecture the schedule is then disseminated to each DS-T station to designate when and what type of service is to be performed by that station. From this high level description of the service, the DS-T proceeds to schedule station specific activities in order to provide the desired services. These activities consist of track script generation, and execution of the track script for each track.

The script generator (SG) is where the majority of the control autonomy is provided. The SG uses Artificial Intelligence planning techniques to perform a complex software module reconfiguration process. This process consists of piecing together numerous highly interdependent smaller control scripts in order to produce a single script to control the operations of the DS-T station.

The core engine used in the SG is the Automated Scheduling and Planning ENvironment (ASPEN). The ASPEN system is a reusable, configurable, generic planning/scheduling application framework that can be tailored to specific domains to create conflict-free plans or schedules. It has a number of useful features including an expressive modeling language, a constraint management system for representing and maintaining antenna operability and/or resource constraints, a temporal reasoning system and a graphical interface for visualizing plans and states. ASPEN has been adapted to input antenna-tracking goals and automatically produce the required command sequence necessary to create the requested link.

The control script produced by the SG: sets up the track by configuring the station during pre-track; provide the track service requested by commanding the antenna and sub-systems to acquire and maintain lock on the signal throughout mode changes; and cleanup and shutdown the station at the completion of the track.

The original goal of the DS-T task was to build an autonomous control system for a deep space communications station. This system had to meet the following criteria: schedule driven with a high level service request interface; an automated scheduling component for initial scheduling and rescheduling; provide script guided control; ability to generate predicts or use provided predicts; automatically configure pre-pass; utilization of COTS components wherever feasible; operations based on defined but expandable set of services; autonomous error recovery for a defined class of problems; post pass data delivery; and treat ground terminal as a network computer with an RF peripheral.

One of the most important points was the idea of a ground station looking just like a network computer to a user, operator, or mission. This is best demonstrated by an operational scenario. To provide service a user need only login to the DS-T workstation and submit a service request to the scheduling system, or FTP a schedule and service request to a particular file system location. From either of these inputs DS-T would detect the existence of a track/service schedule and proceed to schedule station specific task and configure the station to provide the service and finally when the time comes the track would begin without further user interaction.

As mentioned above, the station reacts to a service request derived schedule generated by an automated scheduling system. It is through the reaction to this schedule that the dynamic track specific control scripts are generated. It is through the execution of these control scripts that the autonomous operations of the station takes place.

In May 1998, the DS-T prototype first demonstrated automated downlink capability of single isolated tracks for the Mars Global Surveyor (MGS) spacecraft. Since May, many multi-day demonstrations have taken place including a six day unat-

tended demonstration. During these demonstrations, a service request for downlink services. a track sequence of events. and spacecraft ephemeris were used to automatically downlink data from the MGS spacecraft.

This paper will further describe an architecture for an autonomous deep space tracking station, DS-T. This DS-T station automates routine operations such as: scheduling and resource allocation, antenna and receiver predict generation, track procedure generation from service requests, and closed loop control and error recovery for the station subsystems. This architecture has been successfully demonstrated through a set of DS-T technology demonstrations. The paper will then go into detail on how artificial intelligence planning and scheduling have been applied to the DS-T in order to provide autonomous control ground station operations. The paper will then discuss the ongoing work, building on the success and knowledge gained from the DS-T automation software, to provide dynamic commanding. This approach will integrate the planning aspect of control with the monitor and control component. By integrating these two aspects of automation, a monitor and control system is able to provide more thorough error recovery with reduced reaction time. This work is being done on a system called CLEaR (Closed Loop Error Recovery). The CLEaR system is also being integrated with a fault detection system (FDIR – Fault Detection, Isolation and Recovery) to provide intelligent analysis of monitor data. Once the intelligent analysis is performed CLEaR will reason about the diagnostics and provide an intelligent response.